arXiv:1703.01009v2 [cs.DS] 6 Mar 2017

Optimal Time and Space Construction of Suffix Arrays and LCP Arrays for Integer Alphabets

Keisuke Goto[1]

[1] *Fujitsu Laboratories Ltd., Kawasaki, Japan.,* goto.keisuke@jp.fujitsu.com

Abstract

Suffix arrays and LCP arrays are one of the most fundamental data structures widely used for various kinds of string processing. Many problems can be solved efficiently by using suffix arrays, or a pair of suffix arrays and LCP arrays. In this paper, we consider two problems for a string of length N, the characters of which are represented as integers in $[1, \ldots, \sigma]$ for $1 \leq \sigma \leq N$; the string contains σ distinct characters, (1) construction of the suffix array and (2) simultaneous construction of both the suffix array and the LCP array. In the word RAM model, we propose algorithms to solve both the problems in $O(N)$ time using $O(1)$ extra words, which are optimal in time and space. Extra words mean the required space except for the space of the input string and output suffix array and LCP array. Our contribution improves the previous most efficient algorithm that runs in $O(N)$ time using $\sigma + O(1)$ extra words for the suffix array construction proposed by [Nong, TOIS 2013], and it improves the previous most efficient solution that runs in $O(N)$ time using $\sigma + O(1)$ extra words for both suffix array and LCP array construction using the combination of [Nong, TOIS 2013] and [Manzini, SWAT 2004].

Another optimal time and space algorithm to construct the suffix array was proposed by [Li et al., arXiv 2016] very recently and independently. Our algorithm is simpler than theirs, and it allows us to solve the second problem in optimal time and space.

1 Introduction

Suffix arrays [17] were proposed as a space efficient alternative to suffix trees, which are one of the most fundamental and powerful tools used for various kinds of string processing. Suffix arrays with LCP arrays [17], which are auxiliary data structures of suffix arrays, are sometimes called enhanced suffix arrays [1], and they can simulate various operations of suffix trees. Suffix arrays or enhanced suffix arrays can be used for efficiently solving problems in various research areas, such as pattern matching [17,19], genome analysis [1,15], text compressions [3,5, 8], and data mining [6,9]. In these applications, one of the main computational bottlenecks is the construction of suffix arrays and LCP arrays.

In this paper, we consider two problems that are for a given string: (1) the suffix array construction and (2) the suffix array and LCP array construction. Optimal time construction algorithms have already been proposed for suffix

arrays [11, 13, 14], and for LCP arrays [12], but they are not space optimal. Moreover, optimal time and space construction algorithms for both problems have not been proposed for the last two decades. We herein propose optimal time and space algorithms for both problems. We assume that an input string of length N consists of an integer alphabet of $[1, \ldots, \sigma]$ for $1 \leq \sigma \leq N$ and contains σ distinct characters. It is the same assumption that appears in [2]. Our computational model is the word RAM model with word size $w = \lceil \log N \rceil$, and basic arithmetic and bit operations take constant time.

Suffix arrays construction. The first linear time (optimal time) algorithms to construct suffix arrays have been proposed simultaneously by several authors [11, 13, 14], and they require at least N extra words, where extra words mean the required space except the space of the input string and output array. Nong proposed the most space efficient and linear time algorithm that requires $\sigma + O(1)$ extra words [20], but it still requires about N extra words in the worst case because σ can be N. Nong's algorithm is based on the induced sorting framework [7, 14, 20, 21] that splits all suffixes into L- and S-suffixes, sorts either of which first, and then sorts the other. We propose an *in-place* linear time algorithm based on the framework, where in-place means that the algorithm runs on optimal space, namely $O(1)$ extra words. The induced sorting framework maintains σ pointers that indicate positions in the suffix array where suffixes are to be inserted. Our algorithm runs in almost the same way as the previous ones [7, 14, 20, 21], but it stores these pointers in the space of the output suffix array and requires only $O(1)$ extra words. Therefore, our algorithm runs in $O(N)$ time and in-place.

The induced sorting framework sorts a subset of suffixes recursively and must store $O(1)$ integers at each recursive step, so it requires $O(\log N)$ words in total if it stores naively. However, this space is omitted in the analysis of previous research [14, 20, 21] based on the framework. To the best of our knowledge, only Franceschini and Muthukrishnan [7] addressed this, and they proposed an algorithm to store the values in-place. Each value of $\lceil \log N \rceil$ bits can be read and written in $O(\log N)$ time in their algorithm because it is stored into the least significant bits (LSBs) of consecutive $\lceil \log N \rceil$ elements of the output suffix array, where the LSBs are unused in the whole of recursive steps. For ease of implementation and for reducing the time complexity by constant factors, we propose a simpler algorithm that stores the values in-place and that takes $O(1)$ time to read and write a value.

Suffix arrays and LCP arrays construction. Most previous research focused on a setting that computes LCP arrays from a given string and its suffix array. Kasai et al. proposed the first linear time (optimal time) algorithm to construct LCP arrays from a given string and its suffix array in $N + O(1)$ extra words [12]. Thereafter, Manzini [18] reduced the workspace by using a data structure that is similar to the one used in Kasai et al.'s algorithm, and they rewrote it to the LCP array. The rewriting process runs in-place, but the construction of the data structure requires $\sigma + O(1)$ extra words, so the algorithm runs in $O(N)$ time using $\sigma + O(1)$ words in total. The second problem can be solved in the same time and space complexities by computing the suffix array using Nong's algorithm [20] and then by computing the LCP array using Manzini's algorithm [18]. We propose an algorithm to construct the data structure used in Manzini's algorithm in $O(N)$ time and in-place with constructing the suffix array in-place using a modified version of our suffix array

construction algorithm. After that, we obtain the suffix array and the data structure, and then transform it to LCP arrays as in Manzini's way in $O(N)$ time and in-place. In total, we obtain the suffix array and LCP array in $O(N)$ time and in-place.

Very recently, an in-place suffix array construction algorithm that runs in $O(N)$ time was proposed by Li et al. [16]. Their algorithm is also based on the induced sorting framework, but the details are quite different from ours.

1. Their algorithm must sort the fewer of L- or S-suffixes first, and then sort the other. Our algorithm is more flexible in that we can choose which of them to sort first. This is the most important difference because this property is required for our optimal time and space algorithm for suffix arrays and LCP arrays construction.

2. Their algorithm stores pointers with Jacobson's succinct data structure [10], which uses tricky word RAM techniques, thereby making the implementation complex. Our algorithm stores pointers in a normal integer array, and it is much simpler than their algorithm.

3. The space of the integers required for the recursion is omitted in their analysis, which is $O(\log N)$ words if it stores naively. We address this, and propose a simple algorithm to store the values in-place.

4. They assume that an input string consists of integers $[1, \ldots, \sigma]$ for $1 \leq \sigma \leq N$. We assume that an input string contains σ distinct characters in addition to their assumption.

This paper is organized as follows. Section 2 introduces some notations and definitions. Section 3 explains the induced sorting framework on which our algorithms are based. Section 4 and Section 5 propose optimal time and space algorithms to construct the suffix array and both the suffix array and LCP array, respectively.

2 Preliminaries

Let Σ be an integer alphabet, the elements of which are $[1, \ldots, \sigma]$ for an integer $\sigma \geq 1$. An element of Σ^* is called a string. The length of a string $\mathbf{T}$ is denoted by $|\mathbf{T}|$. The empty string ϵ is the string of length 0, namely $|\epsilon| = 0$. For a string $\mathbf{T} = xyz$, x, y and z are called a *prefix*, *substring*, and *suffix* of $\mathbf{T}$, respectively. For a string $\mathbf{T}$ of length N, the i-th character of a string $\mathbf{T}$ is denoted by $\mathbf{T}[i]$ for $1 \leq i \leq N$, and the substring of $\mathbf{T}$ that begins at position i and that ends at position j is denoted by $\mathbf{T}[i \ldots j]$ for $1 \leq i \leq j \leq N$. For convenience, we assume that $\mathbf{T}[N] = \$$, where $\$$ is a special character lexicographically smaller than any characters in the string $\mathbf{T}[1 \ldots N-1]$. In this paper, we also assume $\sigma \leq |\mathbf{T}|$ and $\mathbf{T}$ contains distinct σ characters.

A suffix starting at i and its first character are denoted by $\mathbf{T}_i$ and c_i, respectively, and i is called a pointer of $\mathbf{T}_i$. If no confusion occurs, we sometimes use $\mathbf{T}_i$ as a pointer i. For $1 \leq i \leq N$, $\mathbf{T}_i$ is called a small type suffix (S-suffix) if $i = N$ or $\mathbf{T}_i$ is lexicographically smaller than $\mathbf{T}_{i+1}$, and a large type suffix (L-suffix) otherwise. An S-suffix/L-suffix $\mathbf{T}_i$ is also called a left-most-S-suffix/left-most-L-suffix (LMS-suffix/LML-suffix) if $i > 1$ and $\mathbf{T}_{i-1}$ is an L-suffix/S-suffix. For

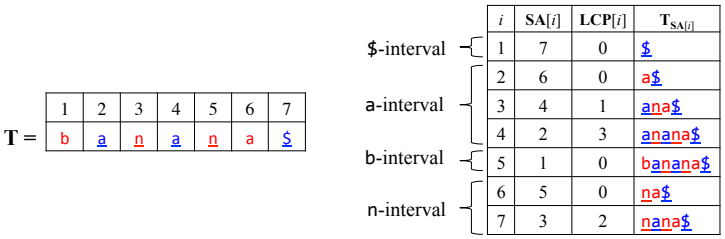

Figure 1: Suffix array and LCP array of $\mathbf{T}$ = banana\$. $\mathbf{T}[i]$ is colored red if $\mathbf{T}_i$ is an L-suffix, and blue otherwise. Moreover, $\mathbf{T}[i]$ is underlined if $\mathbf{T}_i$ is an LML- or LMS-suffix.

$i < N$, $\mathbf{T}'_i$ denotes the substring $\mathbf{T}[i \dots j]$, where $j > i$ is the left most position such that $\mathbf{T}_j$ is an LMS-suffix. Each type of $\mathbf{T}'_i$ is equal to that of $\mathbf{T}_i$, and $\mathbf{T}'_i$ is referred to as an L-, S-, LML-, or LMS-substring according to its type. Note that for $1 \leq i < N$, $\mathbf{T}_i$ is an L-suffix if $c_i > c_{i+1}$, or $c_i = c_{i+1}$ and $\mathbf{T}_{i+1}$ is an L-suffix, and $\mathbf{T}_i$ is an S-suffix otherwise. From this property, each type of suffix can be obtained by a right-to-left scan on $\mathbf{T}$ by comparing the first characters of adjacent suffixes in $O(N)$ time. $suf(all)$ denotes the set of all suffixes of $\mathbf{T}$, and also $suf(L)$, $suf(S)$, $suf(LML)$, and $suf(LMS)$ denote the set of all L-, S-, LML-, and LMS-suffixes of $\mathbf{T}$, respectively. The size of a set M is denoted by N_M. Note that either of $N_{suf(L)}$ or $N_{suf(S)}$ must be less than or equal to $N/2$ because all of the suffixes belong to either one. Moreover, $N_{suf(LMS)}$ must be less than or equal to the smallest of $N_{suf(L)}$ and $N_{suf(S)}$. For a subset M of $suf(all)$, let $\mathbf{T}_j$ be the i-th smallest suffix in M and let the rank of $\mathbf{T}_j$ be denoted by $rank_M(j) = i$, respectively. When the context is clear, we denote $rank_{suf(all)}$ by $rank$.

For a subset M of $suf(all)$, the suffix array $\mathbf{SA}_M$ of length N_M is an integer array that stores all pointers of M such that corresponding suffixes are lexicographically sorted. More precisely, for all suffixes $\mathbf{T}_i$ in M, $\mathbf{SA}_M[rank_M(i)] = i$. When the context is clear, we denote $\mathbf{SA}_{suf(all)}$ by $\mathbf{SA}$. For each character c, the maximum interval in which the first characters of suffixes are equal to c in $\mathbf{SA}$ is called the c-interval. Because L- and S-suffixes are respectively larger and smaller than their succeeding suffix, for any character c, L-suffixes that start with c are always located before S-suffixes starting with c in $\mathbf{SA}$.

The LCP array is an auxiliary array of $\mathbf{SA}$ such that $\mathbf{LCP}[i]$ contains the length of the longest common prefix of $\mathbf{T}_{\mathbf{SA}[i]}$ and $\mathbf{T}_{\mathbf{SA}[i-1]}$ for $1 < i \leq N$, and $\mathbf{LCP}[1] = 0$. See Figure 1.

3 Induced Sorting Framework

Our algorithm is based on the induced sorting framework [7, 14, 20, 21], so in this section, we explain the algorithm proposed by [21], which runs in $O(N)$ time using $N + N/\log N + O(1)$ extra words, as an example of the framework. The key point of the framework is to sort a subset of suffixes once and then sort another subset of suffixes from the sorted subset. From this perspective, we say that the sorting of latter suffixes is induced from the former suffixes. Let $\mathbf{T}^0$ be $\mathbf{T}$ and let $\mathbf{T}^{i+1}$ be a string whose length is the number of LMS-suffixes of

$\mathbf{T}^i$ such that $\mathbf{T}^{i+1}[j] = k$, where the j-th LMS-substring from the left in $\mathbf{T}^i$ is the k-th lexicographically smallest LMS-substring. The algorithm recursively computes the suffix array $\mathbf{SA}^i$ of the string $\mathbf{T}^i$ at each recursive step i.

3.1 Overview of the Algorithm

Below is the overview of the algorithm to compute $\mathbf{SA}^i$ for $\mathbf{T}^i$ at a recursive step i. Note that all suffixes and substrings that appear in the overview indicate those of $\mathbf{T}^i$.

1. Sort all LMS-substrings.
   (a) Sort all LMS-substrings by their first characters.
   (b) Sort all L-substrings from LMS-substrings sorted by their first characters.
   (c) Sort all S-substrings from sorted L-substrings.
2. Sort all LMS-suffixes.
   (a) Check whether or not all LMS-substrings are unique, and if so, do nothing and go to Step 3.
   (b) Create a new string $\mathbf{T}^{i+1}$.
   (c) Compute $\mathbf{SA}^{i+1}$ of $\mathbf{T}^{i+1}$ recursively.
   (d) Convert $\mathbf{SA}^{i+1}$ to the suffix array of LMS-suffixes.
3. Sort all suffixes.
   (a) Do preprocessing for Step 3b.
   (b) Sort all L-suffixes from sorted LMS-suffixes.
   (c) Sort all S-suffixes from sorted L-suffixes.

3.2 Details of the Algorithm

We herein explain the details of the algorithm. Because the essence of the algorithm is the part of sorting all suffixes from sorted LMS-suffixes in Step 3, we first explain Step 3 and then explain Step 1 and 2. We consider only the case to compute $\mathbf{SA}^0 = \mathbf{SA}$ of $\mathbf{T}^0 = \mathbf{T}$ at the recursive step 0, but we can also compute $\mathbf{SA}^i$ of $\mathbf{T}^i$ similarly at each recursive step i.

The algorithm requires three arrays. $\mathbf{A}$ is an integer array of length N to be $\mathbf{SA}$ in the end of the algorithm. $\mathbf{type}$ is a boolean array of length N, which indicates the type of $\mathbf{T}_j$ such that $\mathbf{type}[j] = L$ if $\mathbf{T}_j$ is an L-suffix and $\mathbf{type}[j] = S$ otherwise. $\mathbf{type}$ can be computed in $O(N)$ time by a right to left scan on $\mathbf{T}$ by comparing the first characters of the current suffix and its succeeding suffix. $\mathbf{LE}/\mathbf{RE}$ is an integer array of length σ such that $\mathbf{LE}[c]/\mathbf{RE}[c]$ indicates the next inserting position of a suffix starting with a character c in $\mathbf{A}$. $\mathbf{LE}[c]/\mathbf{RE}[c]$ is initially set to the head/tail of the c-interval, and it is managed to indicate the left-most/right-most empty position of the c-interval at any step of the algorithm. Let C_c be the number of suffixes starting with c. $\mathbf{LE}$ can be initialized in $O(N)$ time as follows: (1) computes $\mathbf{LE}[c] = C_c$ for all characters c by counting $\mathbf{T}[i]$ from right to left using $\mathbf{LE}[c_i]$ as a counter, (2) computes

$\mathbf{LE}[c] = 1 + \sum_{c'<c} C_{c'}$ by accumulating $\mathbf{LE}[c] = C_c$ lexicographically. Similarly, $\mathbf{RE}$ can also be computed in $O(N)$ time.

We assume that $\mathbf{LE}$ is initialized in the beginning of Step 1b and Step 3b, and $\mathbf{RE}$ is also in the beginning of Step 1a, Step 1c, Step 3a, and Step 3c.

Step 3a: Because of the result of Step 2, $\mathbf{SA}_{suf(LMS)}$ is stored in $\mathbf{A}[1 \dots N_{suf(LMS)}]$ and $\mathbf{A}[N_{suf(LMS)} \dots N]$ is filled with empty. By a right to left scan on $\mathbf{SA}_{suf(LMS)}$, we move each $\mathbf{SA}_{suf(LMS)}[i] = \mathbf{T}_j$ into $\mathbf{A}[\mathbf{RE}[c_j]]$, which is the right most empty position of the c_j-interval and decrease $\mathbf{RE}[c_j]$ by one to update the right most empty position of the c_j-interval.

Step 3b: By a left to right scan on $\mathbf{A}$, we read all L- and LMS-suffixes $\mathbf{A}[i] = \mathbf{T}_j$ in lexicographic order, and store $\mathbf{T}_{j-1}$ into $\mathbf{A}[\mathbf{LE}[c_{j-1}]]$ and increase $\mathbf{LE}[c_{j-1}]$ by one if $\mathbf{T}_{j-1}$ is an L-suffix. The type of $\mathbf{T}_{j-1}$ is judged by $\mathbf{type}[j-1]$.

Step 3c: This step runs almost the same way as Step 3b. By a right to left scan on $\mathbf{A}$, all L- and S-suffixes $\mathbf{A}[i] = \mathbf{T}_j$ are read in reverse lexicographic order and $\mathbf{T}_{j-1}$ is stored into $\mathbf{A}[\mathbf{RE}[c_{j-1}]]$ and decrease $\mathbf{RE}[c_{j-1}]$ by one if $\mathbf{T}_{j-1}$ is an S-suffix.

Step 3a, 3b, and 3c run in $O(N)$ time because each step scans $\mathbf{A}$ only one time, and any operations take constant time per access. From Lemma 1, all induced L- and S-suffixes $\mathbf{T}_{j-1}$ are stored in $\mathbf{A}[rank(j-1)]$, so $\mathbf{A} = \mathbf{SA}$ is obtained in the end of Step 3.

Lemma 1 ([21]). *When an L-suffix/S-suffix $\mathbf{T}_{j-1}$ is being induced in Step 3b/Step 3c, $\mathbf{LE}[j-1]$/$\mathbf{RE}[j-1]$ indicates $rank(j-1)$.*

Because Lemma 1 is also important for our algorithm, we left the essence of the proof for Step 3b here. The same things can be said for Step 3c. The proof comes from two important observations. The first observation is that all L-suffixes $\mathbf{T}_{j-1}$ are induced in the c_{j-1}-interval from the properties: (1) $rank(j-1) > rank(j)$, (2) $c_{j-1} \geq c_j$, and (3) $\mathbf{LE}[c_{j-1}]$ always indicate the left most empty position of the c_{j-1}-interval. This property implies $\mathbf{LE}[c_{j-1}] > i$, so $\mathbf{T}_{j-1}$ must be read after reading $\mathbf{T}_j$, and we must not miss any L- and LMS-suffixes on the scan on $\mathbf{A}$. The second observation is that $\mathbf{T}_{j-1}$ must be largest among already stored L-suffixes in the c_{j-1}-interval. It trivially holds when $\mathbf{T}_{j-1}$ is to be stored in the head of the c_{j-1}-intervals. Without a loss of generality, we assume $\mathbf{A}[i] = \mathbf{T}_j$ is the first element where the L-suffix $\mathbf{T}_{j-1}$ is not larger than already stored k suffixes in the c_{j-1}-interval, where k must be larger than 0. From the assumption and the first observation, $\mathbf{T}_j$ and all L- and LMS-suffixes smaller than $\mathbf{T}_j$ are stored in $\mathbf{A}[1 \dots i]$ lexicographically. Let $\mathbf{T}_{p-1}$ be a suffix larger than $\mathbf{T}_{j-1}$, which is already stored in the $\mathbf{T}_{j-1}$-interval and induced from $\mathbf{T}_p$ stored in $\mathbf{A}[1 \dots i-1]$. Because $\mathbf{T}_{p-1}$ and $\mathbf{T}_{j-1}$ starting with the same character c_{j-1} and $\mathbf{T}_{p-1}$ is larger than $\mathbf{T}_{j-1}$, $\mathbf{T}_p$ is also larger than $\mathbf{T}_j$. However, this contradicts the assumption that all suffixes are stored in $\mathbf{A}[1 \dots i]$ lexicographically. From the first and second observation, $\mathbf{LE}[c_{j-1}]$ always indicates $rank(j-1)$ in Step 3b.

Step 1a: By a right to left scan on $\mathbf{T}$, we store all LMS-substrings $\mathbf{T}'_i$ into $\mathbf{A}[\mathbf{RE}[c_i]]$ and decrease $\mathbf{RE}[c_i]$ by one.

Step 1b and 1c: Both are almost the same as Step 3b and 3c, respectively. The important difference is that all substrings other than LMS-substrings are deleted in the end of Step 1. In Step 1b, after inserting an L-substring $\mathbf{T}'_{j-1}$ induced from an L- or LMS-substring $\mathbf{A}[i] = \mathbf{T}'_j$ into $\mathbf{A}[\mathbf{LE}[c_{j-1}]]$, $\mathbf{A}[i]$ is set to empty. After that, $\mathbf{A}$ contains only sorted LML-substrings. In Step 1c, after

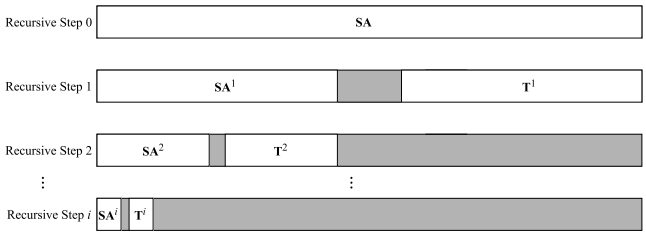

Figure 2: Inside transition of $\mathbf{A}$ at each recursive step i. The space of $\mathbf{SA}^i$ at a recursive step i is used to store $\mathbf{SA}^{i+1}$ and $\mathbf{T}^{i+1}$ for the next recursive step $i+1$.

inserting an S-substring $\mathbf{T}'_{j-1}$ induced from an LML- or S-substring $\mathbf{A}[i] = \mathbf{T}'_j$ into $\mathbf{A}[\mathbf{RE}[c_{j-1}]]$, $\mathbf{A}[i]$ is set to empty. In the end of Step 1, $\mathbf{A}$ contains only sorted LMS-substrings.

Step 2: We only explain the essence of Step 2 because this step runs in $O(N)$ time and in-place, except for the space required for recursion in Step 2c, and our algorithm runs in a similar way for this step. If LMS-substrings are all unique, pointers of sorted LMS-substrings are equal to those of LMS-suffixes, so go to Step 3 immediately. Otherwise, we create $\mathbf{T}^1$ such that each LMS-substring in $\mathbf{T}$ is replaced by its rank among LMS-substrings, and store $\mathbf{T}^1$ in $\mathbf{A}[N - |\mathbf{T}^1| + 1 \dots N]$. We compute the suffix array $\mathbf{SA}^1$ of $\mathbf{T}^1$ recursively using $\mathbf{A}[1 \dots |\mathbf{T}^1|]$ as a working array to be $\mathbf{SA}^1$ for the next recursive step. See Figure 2. Here, a pointer of $\mathbf{T}^1[i]$ is related by the pointer of i-th LMS-substring from the left in $\mathbf{T}$. After that, we convert all pointers of $\mathbf{SA}^1$ into the corresponding ones of $\mathbf{SA}_{suf(LMS)}$. Because $|\mathbf{T}^1| = N_{suf(LMS)}$ is less than or equal to $N/2$, $\mathbf{SA}^1$ and $\mathbf{T}^1$ do not overlap in $\mathbf{A}$.

For each recursive step i, the algorithm takes $O(|\mathbf{T}^i|)$ time. Because $|\mathbf{T}^{i+1}|$ must be smaller than half of $|\mathbf{T}^i|$, the algorithm takes $O(N)$ time in total. Moreover, it requires $N + N/\log N + O(1)$ extra words. The first and second factors are for $\mathbf{LE}/\mathbf{RE}$ and **type**, respectively. See [21] for more details.

4 Optimal Time and Space Construction of Suffix Arrays

We propose a new suffix array construction algorithm based on the induced sorting framework in $O(N)$ time and in-place. The space bottleneck of the previous algorithm [21] is the space of $\mathbf{LE}/\mathbf{RE}$ and **type**. We reduce the space by storing $\mathbf{LE}/\mathbf{RE}$ in the space of $\mathbf{A}$, deleting **type** and getting the types of induced suffixes in another way.

As seen in Section 3, the essence of the induced sorting framework is the part of sorting L-suffixes. First, we explain the idea and algorithm to sort all L-suffixes from sorted LMS-suffixes in $O(N)$ time and in-place in Section 4.1 and Section 4.2. Second, we show the same idea can be applied to sort all S-suffixes and LMS-substrings in Section 4.3 and Section 4.4, respectively. Third, we propose a new method to store $O(\log N)$ integers required for the recursion in-place in Section 4.5. We omit the details of sorting LMS-suffixes from sorted LMS-substrings because our algorithm computes that in $O(N)$ time and in-place

in the same way as the previous algorithm [21]. The important property of our algorithm is that optionally we can sort LML-suffixes first, S-suffixes next, and then L-suffixes by switching the roles of L-suffixes and S-suffixes.

In the rest of the section, we assume the situation to compute $\mathbf{SA}^0 = \mathbf{SA}$ of $\mathbf{T}^0 = \mathbf{T}$, but similarly we can also compute $\mathbf{SA}^i$ of $\mathbf{T}^i$ at each recursive step i.

4.1 Main Idea and Preliminaries for Sorting L-suffixes

We give the main idea, some definitions, and notations for our algorithm.

Our main idea for reducing the space of **LE** is to store it somewhere in **A**. Intuitively, this idea does not work because elements of **LE** may be overwritten by induced L-suffixes even though they may be required in the future. Moreover, we may not be able to even store $\mathbf{SA}_{suf(LMS)}$ and **LE** in **A** at the same time before sorting L-suffixes because the sum of their sizes may be greater than N. We can avoid this problem using two key tricks. The first trick is to store some of the L- and LMS-suffixes $\mathbf{T}_i$ in $\mathbf{LE}[c_i]$ when all L-suffixes starting with c_i have been induced and when $\mathbf{LE}[c_i]$ is no longer used in the future. The second trick is to store sorted L- and LMS-suffixes in three internal sub arrays in **A**. As a by-product of the second trick, judging the type of a suffix can be done without **type** in our algorithm. Therefore, we can reduce the space of **LE** and **type**, and the algorithm runs in-place.

Let $suf(LMSx)$ be the set of all LMS-suffixes that are the smallest LMS-suffixes starting with each character c, let $suf(\overline{LMSx})$ be the set of the other LMS-suffixes, let $suf(LMSy)$ be the set of all LMS-suffixes that are the smallest LMS-suffixes starting with each character c such that there is no L-suffix starting with c, let $suf(\overline{LMSy})$ be the other LMS-suffixes, let $suf(Lx)$ be the set of all L-suffixes that are the largest L-suffixes starting with each character c, and let $suf(\overline{Lx})$ be the other L-suffixes. $suf(LMSy) \cup suf(Lx)$ can be regarded as a set of suffixes that includes, for each c-interval, the largest L-suffix if it exists and the smallest LMS-suffix (which also may not exist) otherwise. Therefore, it contains at most σ suffixes, and the first characters are all different. Note that a suffix starting with a character c is not contained in the set if any S-suffix starts with c but no L- and LMS-suffixes start with c. For an array $\mathbf{Y}$ of length σ and a set of suffixes M of which the first characters are distinct, $\mathbf{Y}_M$ denotes $\mathbf{Y}$ *overwritten* by M, meaning that for a character c, $\mathbf{Y}_M[c] = \mathbf{T}_i$ if $\mathbf{T}_i$ starting with c exists in M, and $\mathbf{Y}_M[c] = \mathbf{Y}[c]$ otherwise.

In our algorithm, **LE** is redefined to manage the next inserting position of a suffix in $\mathbf{SA}_{suf(\overline{Lx})}$ not **SA** because **LE** is used for constructing $\mathbf{SA}_{suf(\overline{Lx})}$. Let CL_c be the number of L-suffixes starting with c. Note that no c-interval is in $\mathbf{SA}_{suf(\overline{Lx})}$ such that $CL_c \leq 1$. For a character c, $\mathbf{LE}[c]$ is initially set to the head of the c-interval if $CL_c > 1$, the left most position i such that $c_{\mathbf{SA}_{suf(\overline{Lx})}[i]} > c$ if $CL_c = 1$, and it is empty otherwise. The algorithm uses three internal sub arrays of **A**, $\mathbf{X} = \mathbf{A}[1 \dots N_{suf(L)}]$, $\mathbf{Y} = \mathbf{A}[N - N_{suf(\overline{LMSy})} - \sigma + 1 \dots N - N_{suf(\overline{LMSy})}]$, and $\mathbf{Z} = \mathbf{A}[N - N_{suf(LMS)} + 1 \dots N]$. $\mathbf{X}$ is split into $\mathbf{X}_1 = \mathbf{A}[1 \dots N_{suf(\overline{Lx})}]$ and $\mathbf{X}_2 = \mathbf{A}[N_{suf(\overline{Lx})} + 1 \dots N_{suf(L)}]$, $\mathbf{Z}$ is split into $\mathbf{Z}_1 = \mathbf{A}[N - N_{suf(LMS)} + 1 \dots N - N_{suf(\overline{LMSx})}]$ and $\mathbf{Z}_2 = \mathbf{A}[N - N_{suf(\overline{LMSx})} + 1 \dots N]$, and $\mathbf{Z}$ has another sub array $\mathbf{Z}_3 = \mathbf{A}[N - N_{suf(\overline{LMSy})} + 1 \dots N]$. Each size of sets $suf(LMSx)$, $suf(\overline{LMSx})$, $suf(LMSy)$, $suf(\overline{LMSy})$, $suf(Lx)$, and $suf(\overline{Lx})$ can be easily computed in either of Transition 2, 4, and 6. For ease of explanation, we assume their sizes are

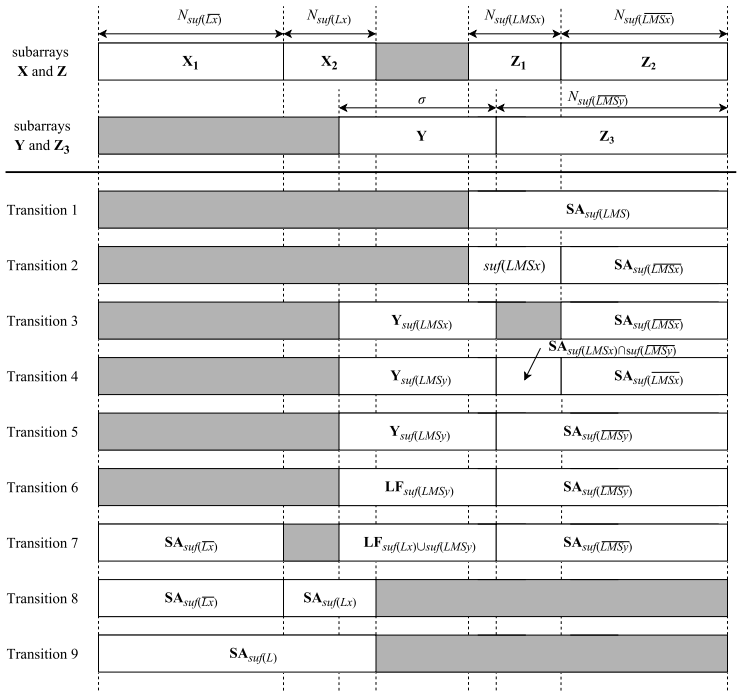

Figure 3: Inside transition of $\mathbf{A}$ in the progress of computing $\mathbf{SA}_{suf(L)}$ from $\mathbf{SA}_{suf(LMS)}$. Gray space denotes empty space.

already computed before sorting L-suffixes, and any index of $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ can be accessed in $O(1)$ time.

$\mathbf{Y}$ is used for storing suffixes $\mathbf{T}_j$ of $suf(Lx) \cup suf(LMSy)$ in $\mathbf{Y}[c_j]$. Each type of a suffix $\mathbf{Y}[c_j] = \mathbf{T}_j$ is judged by a boolean array $\mathbf{typeY}$ of length σ such that $\mathbf{typeY}[c_j] = S$ if $\mathbf{T}_j \in suf(LMSy)$, and $\mathbf{typeY}[c_j] = L$ if $\mathbf{T}_j \in suf(Lx)$. Though $\mathbf{typeY}$ requires σ bits if it is stored naively, it requires no extra space because we embed $\mathbf{typeY}[c_j]$ into the least significant bit (LSB) of $\mathbf{Y}[c_j]$. For this purpose, an L- or LMS-suffix $\mathbf{T}_j$ is represented by the compact representation $\lfloor j/2 \rfloor$ so that the LSB can be used for $\mathbf{typeY}[c_j]$. When reading a compact representation k stored in $\mathbf{Y}[c]$, the original suffix $\mathbf{T}_j$ can be recovered in $O(1)$ time. Two important properties are evident; (1) the first character of the original suffix is c because $\mathbf{Y}[c]$ must store a suffix starting with c, (2) the original suffix is either of the largest L-suffix or of the smallest LMS-suffix. Here the candidates of $\mathbf{T}_j$ are $\mathbf{T}_{2k}$ and $\mathbf{T}_{2k+1}$. If $c_{2k} \neq c_{2k+1}$, $\mathbf{T}_j$ is either of the suffixes that starts with c_j, where c_j is obtained from the first property. Otherwise, $\mathbf{T}_j$ is $\mathbf{T}_{2k}$ from the second property. Because it can be computed in $O(1)$ time for both cases, $suf(Lx) \cup suf(LMSy)$ and $\mathbf{typeY}$ can be stored in $\mathbf{Y}$.

4.2 Sort all L-suffixes

We make $\mathbf{SA}_{suf(L)}$ in the head of $\mathbf{A}$ from $\mathbf{SA}_{suf(LMS)}$ stored in the head of $\mathbf{A}$, which is given by the result of sorting LMS-suffixes in the previous step. The internal transitions of $\mathbf{A}$ are shown in Figure 3. Because only $N_{suf(\overline{Lx})} +$

$\sigma + N_{suf(\overline{LMSy})} \leq N$ elements of $\mathbf{A}$ are used in the entire transition of $\mathbf{A}$ and because each transition takes $O(N)$ time, our algorithm runs in $O(N)$ time and in-place.

Each transition is processed as follows.

Transition 1: We shift $\mathbf{SA}_{suf(LMS)}$ stored in the head of $\mathbf{A}$ into $\mathbf{Z}$.

Transition 2: We store unsorted $suf(LMSx)$ in $\mathbf{Z}_1$ and make $\mathbf{SA}_{suf(\overline{LMSx})}$ in $\mathbf{Z}_2$. Let j be the inserting position in $\mathbf{Z}_2$ for $\mathbf{SA}_{suf(\overline{LMSx})}$, which is initially set to $N_{suf(\overline{LMSx})}$. By a right to left scan on $\mathbf{SA}_{suf(LMS)}$, we swap $\mathbf{SA}_{suf(LMS)}[i] = \mathbf{T}_k$ with $\mathbf{Z}_2[j]$ and decrease j by one if $\mathbf{T}_k \in suf(\overline{LMSx})$ and do nothing otherwise. Because the first characters $c_{\mathbf{SA}_{suf(LMS)}[i]}$ and $c_{\mathbf{SA}_{suf(LMS)}[i-1]}$ are the same if $\mathbf{T}_k \in suf(\overline{LMSx})$, whether or not $\mathbf{T}_k$ belongs to $suf(\overline{LMSx})$ can be judged in $O(1)$ time by comparing the first characters. Because we only shift the suffixes of $suf(\overline{LMSx})$ into the tail of $\mathbf{Z}_2$ preserving the sorted order, we obtain unsorted $suf(LMSx)$ in $\mathbf{Z}_1$ and $\mathbf{SA}_{suf(\overline{LMSx})}$ in $\mathbf{Z}_2$.

Transition 3: We make $\mathbf{Y}_{suf(LMSx)}$. By a right to left scan on $\mathbf{Z}_1$, we try to move $\mathbf{Z}_1[i] = \mathbf{T}_{j_1}$ into $\mathbf{Y}[c_{j_1}]$. However, $\mathbf{Y}[c_{j_1}]$ may contain an LMS-suffix $\mathbf{T}_{j_2}$ because $\mathbf{Y}$ may overlap with $\mathbf{Z}_1$. If $\mathbf{Y}[c_{j_1}]$ is empty or $\mathbf{T}_{j_1}$, we simply move $\mathbf{T}_{j_1}$ into $\mathbf{Y}[c_{j_1}]$. Otherwise, we move $\mathbf{T}_{j_1}$ into $\mathbf{Y}[c_{j_1}]$ and then try to move $\mathbf{T}_{j_2}$ into $\mathbf{Y}[c_{j_2}]$. We repeat this procedure until $\mathbf{T}_{j_k}$ is moved into $\mathbf{Y}[c_{j_k}]$, which is empty or $\mathbf{T}_{j_k}$. Because the first characters of $suf(LMSx)$ are all different, the number of insertions is $O(\sigma)$, and the aforementioned process runs in $O(\sigma)$ time.

Transition 4: We make $\mathbf{Y}_{suf(LMSy)}$ and $\mathbf{SA}_{suf(\overline{LMSy})\cap suf(LMSx)}$. We first make $\mathbf{typeY}$. By a single scan on $\mathbf{Y}_{suf(LMSy)}$, we convert all LMS-suffixes to compact representation and initialize $\mathbf{typeY}[c] = S$ for all characters c. In the rest of the section, all L- and LMS-suffixes stored in $\mathbf{Y}$ are represented by compact representation, but ones stored in the other sub arrays are not. For ease of explanation, $\mathbf{Y}[c_j] = \mathbf{T}_j$ denotes a suffix $\mathbf{T}_j$ in both the original and compact representations if no confusion occurs. By a right to left scan on $\mathbf{T}$, we set $\mathbf{typeY}[c_i]$ to L if $\mathbf{T}_i$ is an L-suffix. Let i be the inserting position in $\mathbf{Z}_1$ for $\mathbf{SA}_{suf(\overline{LMSy})}$, which is initially set to $N_{suf(LMSx)}$. By a right to left scan on $\mathbf{Y}_{suf(LMSx)}$, we move $\mathbf{Y}_{suf(LMSx)}[c]$ into $\mathbf{Z}_1[i]$ and decrease i by one if $\mathbf{typeY}[c] = L$. Because we shift only suffixes of $suf(\overline{LMSy}) \cap suf(LMSx)$ into the tail of $\mathbf{Z}_1$ preserving the sorted order, we obtain $\mathbf{Y}_{suf(LMSy)}$ in $\mathbf{Y}$ and $\mathbf{SA}_{suf(\overline{LMSy})\cap suf(LMSx)}$ in the tail of $\mathbf{Z}_1$.

Transition 5: We make $\mathbf{SA}_{suf(\overline{LMSy})}$. Because a suffix $\mathbf{T}_i$ of $suf(\overline{LMSy}) \cap suf(LMSx)$ is smaller than all of the suffixes starting with the same character c_i of $suf(\overline{LMSx})$, $\mathbf{SA}_{suf(\overline{LMSy})}$ can be obtained by the stable merge of $\mathbf{SA}_{suf(\overline{LMSy})\cap suf(LMSx)}$ and $\mathbf{SA}_{suf(\overline{LMSx})}$ considering the first characters as keys. By applying Theorem 1 into $\mathbf{SA}_{suf(\overline{LMSy})\cap suf(LMSx)}$ and $\mathbf{SA}_{suf(\overline{LMSx})}$, we compute $\mathbf{SA}_{suf(\overline{LMSy})}$ in $O(N)$ time and in-place.

Theorem 1 ([4]). *For two integer arrays $\mathbf{A}_1 = \mathbf{A}[1 \dots N_1]$ and $\mathbf{A}_2 = \mathbf{A}[N_1 + 1 \dots N_1 + N_2]$ that are stored in an array $\mathbf{A}$ and that contain sorted integers, respectively, there is an algorithm to merge $\mathbf{A}_1$ and $\mathbf{A}_2$ in $\mathbf{A}[1 \dots N_1 + N_2]$ such that all values are sorted stably in $O(N_1 + N_2)$ time and in-place.*

Transition 6: We make $\mathbf{LE}_{suf(LMSy)}$. By a single scan on $\mathbf{Y}$, $\mathbf{Y}[c]$ is initialized by 0 if $\mathbf{typeY}[c] = L$. By a right to left scan on $\mathbf{T}$, we compute

and store each CL_c into $\mathbf{Y}[c]$ with $\mathbf{Y}[c]$ as a counter. Note that $\mathbf{Y}[c_i]$ storing $\mathbf{T}_i$ of $suf(LMSy)$ can never be overwritten because no L-suffix starting with c_i exists. CL_c can be larger than $2^{\lceil \log N \rceil - 1}$, and if so, the LSB cannot be used for $\mathbf{typeY}[c]$. We avoid this by treating such c as a special character and by storing c and CL_c in two extra words instead of storing them in $\mathbf{Y}[c]$. Because only one special character exists [1], it requires $O(1)$ extra words. By a left to right scan on $\mathbf{Y}$, $CL_c - 1$ is accumulated such that $\mathbf{typeY}[c] = L$, and we set $\mathbf{Y}[c] = \mathbf{LE}[c] = 1 + \sum_{c'<c} \max(0, CL_{c'} - 1)$ if $CL_c > 0$. Otherwise, it means that some S-suffixes start with c, but L- and LMS-suffixes do not start with c. In this case, we set $\mathbf{Y}[c] = \mathbf{LE}[c]$ to empty. Here, we do not use the LSB of $\mathbf{Y}[c]$ after setting $\mathbf{Y}[c] = \mathbf{typeY}[c]$ because $\mathbf{LE}[c]$ can be larger than $2^{\lceil \log N \rceil - 1}$, and if so, the LSB of $\mathbf{Y}[c]$ cannot be used for $\mathbf{typeY}[c]$. However, $\mathbf{Y}[c] = \mathbf{T}_j$ is still represented in compact representation and the LSB is used for $\mathbf{typeY}[c] = S$. $\mathbf{typeY}$ is stored in $\mathbf{Y}$ non-uniformly, but this behavior does not affect the rest of the procedure.

Transition 7: We make $\mathbf{SA}_{suf(\overline{Lx})}$ and $\mathbf{LE}_{suf(Lx) \cup suf(LMSy)}$. The concept is almost the same as that of Section 3, consisting of the following three parts. See Part 3 first because Part 1 and 2 read suffixes stored by Part 3, and their procedures depend on that of Part 3. Note that all suffixes stored in $\mathbf{Y}$ are represented in compact representation, and their types are stored in $\mathbf{typeY}$ at the end of the transition. The correctness is proved in Lemma 2.

Part 1: Read all L- and LMS-suffixes $\mathbf{T}_i$ lexicographically. Suffixes stored in $\mathbf{X}_1$, $\mathbf{Y}$, and $\mathbf{Z}_3$ are sorted in them, respectively. By a left to right scan on $\mathbf{X}_1$, $\mathbf{Y}$, and $\mathbf{Z}_3$, we choose and read the smallest one of suffixes $\mathbf{X}_1[i_X]$, $\mathbf{Y}[i_Y]$, and $\mathbf{Z}_3[i_Z]$, where i_X, i_Y, i_Z are the scanning positions of $\mathbf{X}_1$, $\mathbf{Y}$, and $\mathbf{Z}_3$, respectively. By Part 3, $\mathbf{X}_1[i_X]$ contains either of two values: (1) empty, (2) an L-suffix; $\mathbf{Y}[i_Y]$ contains either of three values: (1) empty, (2) a suffix of $suf(Lx) \cup suf(LMSy)$, (3) $\mathbf{LE}[i_Y]$; and $\mathbf{Z}_3[i_Z]$ contains an LMS-suffix. Let c_Y be i_Y, and let c_X and c_Z be the first character of $\mathbf{X}_1[i_X]$ and $\mathbf{Z}_3[i_Z]$, respectively. If $\mathbf{X}_1[i_X]$ is empty or if an index indicates a position out of the corresponding array, we assume the corresponding characters c_X, c_Y, c_Z are larger characters than σ.

We choose the smallest character c_k of c_X, c_Y, and c_Z, where c_X, c_Y, and c_Z have higher priority in this order if ties are found. If $\mathbf{Y}[i_Y]$ is empty, it implies that none of the L- or LMS-suffixes start with c_Y, so we increase $i_Y = c_Y$ by one and choose the smallest character c_k again. As will be described in the next paragraph, $\mathbf{Y}[i_Y]$ must contain a suffix if c_Y is chosen. We read the suffix corresponding to c_k as $\mathbf{T}_i$ and increase the corresponding scanning position by one. Note that $\mathbf{T}_i$ is the smallest of the three candidates because for suffixes $\mathbf{T}_{j_1}, \mathbf{T}_{j_2}, \mathbf{T}_{j_3}, \mathbf{T}_{j_4}$ of $suf(\overline{Lx})$, $suf(Lx)$, $suf(LMSy)$, $suf(\overline{LMSy})$, respectively, it holds that $\mathbf{T}_{j_1} < \mathbf{T}_{j_2} < \mathbf{T}_{j_3} < \mathbf{T}_{j_4}$ if they all start with the same character, and $\mathbf{T}_i$ is chosen in this priority order.

The aforementioned process does not work when c_Y is chosen, but the largest L-suffix starting with c_Y is stored in $\mathbf{X}_1$ and not in $\mathbf{Y}$. We avoid this by adding extra operations only when reading L-suffix $\mathbf{T}_i$ from $\mathbf{X}_1[i_X]$. After reading $\mathbf{T}_i$ and increasing i_X by one, we check whether or not $\mathbf{X}_1[i_X]$ is empty, and if so, it implies that no L-suffix starting with c_i will be induced in the future and that

[1] If two such special characters exist, $\mathbf{T}$ contains characters, the number of which is larger than N.

a conflict did not occur between the largest suffix starting with c_i and a larger suffix in Part 3; thus, the $\mathbf{T}_i$ belongs to $suf(Lx)$. In this case, we decrease i_X by one, and we move $\mathbf{T}_i$ in $\mathbf{X}_1[i_X]$ into $\mathbf{Y}[i_Y]$, and increase i_Y by one because $\mathbf{Y}[i_Y] = \mathbf{T}_i$ has been already read, where $i_Y = c_Y = c_i$ because $i_Y = c_Y$ is increased one by one and is larger than or equal to c_i. In this way, $\mathbf{Y}[i_Y]$ must store a suffix when c_Y is chosen.

Part 2: Judge whether $\mathbf{T}_{i-1}$ is an L-suffix or not. If $\mathbf{T}_i$ is read from $\mathbf{X}_1$ or $\mathbf{Z}_3$, we know the type of $\mathbf{T}_i$, so the type of $\mathbf{T}_{i-1}$ can be obtained by comparing their first characters. Otherwise, $\mathbf{T}_i$ is read from $\mathbf{Y}$, and we know it belongs to $suf(Lx) \cup suf(LMSy)$. It implies that the first characters of c_i and its preceding character c_{i-1} must be different, so the type of $\mathbf{T}_{i-1}$ can be obtained by comparing their first characters.

Part 3: Store $\mathbf{T}_{i-1}$ if it is an L-suffix. We try to store $\mathbf{T}_{i-1}$ in $\mathbf{X}_1[\mathbf{LE}[c_{i-1}]]$, which is the next inserting position of a suffix starting with c_{i-1}. However, conflicts may occur because we try to store CL_c L-suffixes starting with c into the c-interval of the length $CL_c - 1$ in $\mathbf{X}_1$. Because we try to store the largest suffix starting with c just after the c-interval, which is the head of a c'-interval such that $c' > c$, the suffix starting with the smaller first character of the conflicted two suffixes belongs to $suf(Lx)$ if conflicts occur.

If $\mathbf{X}_1[\mathbf{LE}[c_{i-1}]]$ is empty, we simply store $\mathbf{T}_{i-1}$ there, and increase $\mathbf{LE}[c_{i-1}]$ by one to update the next inserting position. Otherwise, a conflict occurs with $\mathbf{T}_{i-1}$ and $\mathbf{T}_j$ stored in $\mathbf{X}_1[\mathbf{LE}[c_{i-1}]]$. By comparing their first characters, we store the largest one into $\mathbf{X}_1[\mathbf{LE}[c_{i-1}]]$, and store the other into $\mathbf{Y}[\min(c_{i-1}, c_j)]$. If $c_{i-1} > c_j$, increase $\mathbf{LE}[c_{i-1}]$ by one to update the next inserting position.

Lemma 2. *When an L-suffix $\mathbf{T}_{i-1}$ is being induced in Part 3 of Transition 7, $\mathbf{LE}[i-1]$ indicates $rank_{suf(\overline{Lx}) \cup \{\mathbf{T}_{i-1}\}}(i-1)$.*

Proof. If $\mathbf{T}_{i-1}$ belongs to $suf(\overline{Lx})$, $rank_{suf(\overline{Lx}) \cup \{\mathbf{T}_{i-1}\}}(i-1)$ equals $rank_{suf(\overline{Lx})}(i-1)$ and also $\mathbf{LE}[i-1]$, which can be proved in a similar way to Lemma 1. If $\mathbf{T}_{i-1}$ belongs to $suf(Lx)$ and $CL_{c_{i-1}} = 1$, $\mathbf{LE}[i-1]$ indicates $rank_{suf(\overline{Lx}) \cup \{\mathbf{T}_{i-1}\}}(i-1)$ from the definition. Otherwise, the largest suffix $\mathbf{T}_k$ starting with c_{i-1} of $suf(\overline{Lx})$ has already been stored in $\mathbf{A}[rank_{suf(\overline{Lx})}(k)]$, and $\mathbf{LE}[c_{i-1}]$ indicates $1 + rank_{suf(\overline{Lx})}(k)$. Because $\mathbf{T}_k$ is the largest suffix of $suf(\overline{Lx})$ that is smaller than $\mathbf{T}_{i-1}$, $\mathbf{LE}[c_{i-1}]$ indicates $rank_{suf(\overline{Lx}) \cup \{\mathbf{T}_{i-1}\}}(i-1)$. □

Transition 8: We make $\mathbf{SA}_{suf(Lx)}$ in $\mathbf{X}_2$ and initialize the space except $\mathbf{X}$ in $\mathbf{A}$ as empty. Note that $suf(Lx) \cup suf(LMSy)$ are stored in $\mathbf{Y}$ and are represented in compact representation, and $\mathbf{typeY}$ is embedded in the LSBs of $\mathbf{Y}$. Let j be the inserting position in $\mathbf{X}_2$ for $\mathbf{SA}_{suf(Lx)}$, which is initially set to 1. By a left to right scan on $\mathbf{Y}$, we move $\mathbf{Y}[c] = \mathbf{T}_i$ into $\mathbf{X}_2[j]$ if $\mathbf{Y}[c]$ is not empty and $\mathbf{typeY}[c] = L$, and we increase j by one. After that, we obtain $\mathbf{SA}_{suf(Lx)}$ in $\mathbf{X}_2$ and then initialize $\mathbf{A}[N_{suf(L)} \dots N]$ as empty.

Transition 9: We make $\mathbf{SA}_{suf(L)}$. By applying Theorem 1 into $\mathbf{SA}_{suf(Lx)}$ and $\mathbf{SA}_{suf(\overline{Lx})}$ considering the first characters as keys, we compute $\mathbf{SA}_{suf(L)}$ in $O(N)$ time and in-place.

4.3 Sort all S-suffixes

We can sort all S-suffixes in almost the same way of sorting L-suffixes but make **SA** instead of $\mathbf{SA}_{suf(S)}$. The same things can be said by switching the roles of **LE**, L- and LMS-suffixes to **RE**, S- and L-suffixes, respectively. Let $suf(Sx)$ be the smallest suffixes starting with each character c, let $suf(\overline{Sx})$ be the set of the other S-suffixes, let $suf(Ly)$ be the set of the largest L-suffixes starting with each character c such that no S-suffix starts with c, and let $suf(\overline{Ly})$ be the set of the other L-suffixes. We compute $\mathbf{SA}_{suf(\overline{Ly})}$, $\mathbf{RE}_{suf(Ly)\cup suf(Sx)}$ and $\mathbf{SA}_{suf(\overline{Sx})}$ from $\mathbf{SA}_{suf(L)}$ in a similar way of Transition 1-7 in Section 4.2. Note that $\mathbf{RE}_{suf(Ly)\cup suf(Sx)}$ equals $\mathbf{SA}_{suf(Ly)\cup suf(Sx)}$ from the definition. We compute **SA** in $O(N)$ time and in-place by considering the head characters as keys and by applying Theorem 1 into $\mathbf{SA}_{suf(\overline{Ly})}$ and $\mathbf{SA}_{suf(Ly)\cup suf(Sx)}$ and then by applying the result and $\mathbf{SA}_{suf(\overline{Sx})}$.

In this way, sorting all S-suffixes can be done in $O(N)$ time and in-place as in Section 4.2.

4.4 Sort all LMS-substrings

We propose an algorithm to sort LMS-substrings by their first characters in $O(N)$ time and in-place. If we have LMS-substrings sorted by their first characters, we can compute sorted LMS-substrings in $O(N)$ time and in-place in the same way as in the previous algorithm [21] described in Section 3, where L- and S-substrings are sorted in the same way as Section 4.2 and 4.3.

Let $\mathbf{Y} = \mathbf{A}[N - N_{suf(\overline{LMSx})} - \sigma + 1 \dots N - N_{suf(\overline{LMSx})}]$ and $\mathbf{Z} = \mathbf{A}[N - N_{suf(\overline{LMSx})} + 1 \dots N]$. We make **RE** for $\mathbf{SA}_{suf(\overline{LMSx})}$ in $\mathbf{Y}$. By a right to left scan on $\mathbf{T}$, we try to store an LMS-substring $\mathbf{T}'_i$ into $\mathbf{Z}[\mathbf{RE}[c_i]]$. In the same way as in Section 4.2, conflicts may occur, and if so, we store a substring $\mathbf{T}'_k$ starting with the smaller character into $\mathbf{Y}[c_k]$ and a substring starting with the larger character into $\mathbf{Z}[\mathbf{RE}[c_i]]$. After that, LMS-substrings sorted by the first characters are separately stored in $\mathbf{Y}$ and $\mathbf{Z}$. Because $\mathbf{Y}$ stores LMS-substrings sparsely, we shift them into $\mathbf{Y}_1 = \mathbf{A}[N - N_{suf(LMS)} + 1 \dots N - N_{suf(\overline{LMSx})}]$. By applying Theorem 1 into $\mathbf{Y}_1$ and $\mathbf{Z}$ considering the first characters as keys, we obtain $suf(LMS)$ sorted by the first characters in $\mathbf{A}[N - N_{suf(LMS)} + 1 \dots N]$.

4.5 Store $O(\log N)$ Integers for Recursion

The induced sorting framework requires the beginning position of $\mathbf{T}^i$ in $\mathbf{A}$ and its length at each recursive step i. We propose a simple method to store and get each value in $O(1)$ time and in-place.

Let N_i be $|\mathbf{T}^i|$. The key property is that N_{i+1} is at most $\lfloor N_i/2 \rfloor$ because N_{i+1} equals the number of LMS-substrings of $\mathbf{T}^i$. We store $\mathbf{SA}^i$ and $\mathbf{T}^i$ into $\mathbf{A}[1 \dots N_i]$ and $\mathbf{A}[M_i + 1 \dots M_i + N_i]$, where $M_i = \lfloor N/2^i \rfloor$ is the upper bound of N_i. Each M_i, which is the beginning position of $\mathbf{T}^i$ in $\mathbf{A}$, can be computed in $O(1)$ time and in-place by right shifting N by i bits .

Let $\mathbf{B}$ be a boolean array such that $\mathbf{B}[i] = 1$ if $N_i = M_i$, and 0 otherwise. If $\mathbf{B}[i] = 1$, we do not need to store N_i because $N_i = M_i$ and M_i can be obtained in $O(1)$ time and in-place. Otherwise, $\mathbf{A}[N_i + 1 \dots M_i]$ is unused for computing $\mathbf{SA}^i$ of $\mathbf{T}^i$, so we store N_i in $\mathbf{A}[M_i]$. Because the number of recursions is at

most $\lceil \log N \rceil$, $\mathbf{B}$ can be stored in a word. In this way, we can store and get each N_i, which is the length of $\mathbf{T}^i$, in $O(1)$ time and in-place.

5 Optimal Time and Space Construction of Suffix Arrays and LCP Arrays

We propose an algorithm to compute the suffix array and LCP array of a given string $\mathbf{T}$ in $O(N)$ time and in-place. We revisit Manzini's algorithm [18], which constructs the LCP array $\mathbf{LCP}$ from a given string $\mathbf{T}$ and the suffix array $\mathbf{SA}$ in $O(N)$ time using $\sigma + O(1)$ extra words. The algorithm uses an integer array $\mathbf{\Psi}$ of length N where $\mathbf{\Psi}[rank(i)] = rank(i+1)$ for $1 \leq i < N$. Let $\mathbf{B}$ be an integer array of length N to be $\mathbf{LCP}$ at the end of the algorithm. The algorithm consists of two parts. The first part makes $\mathbf{B} = \mathbf{\Psi}$ in $O(N)$ time using $\sigma + O(1)$ extra words. The second part converts $\mathbf{\Psi}$ into $\mathbf{LCP}$ in $O(N)$ time and in-place. Therefore, $\mathbf{LCP}$ can be computed in $O(N)$ time and in-place if $\mathbf{\Psi}$ can be computed in $O(N)$ time and in-place.

We propose a modified version of the algorithm in Section 4, which computes $\mathbf{SA}$ of a given string $\mathbf{T}$ with computing $\mathbf{\Psi}$ in $O(N)$ time and in-place. Therefore, we can obtain $\mathbf{SA}$ and $\mathbf{LCP}$ in $O(N)$ time and in-place. Our new algorithm computes $\mathbf{SA}_{suf(LMS)}$ in the same way as in Section 4, and it computes $\mathbf{SA}$ from $\mathbf{SA}_{suf(LMS)}$ using the following steps, where $suf(Sy)$ is the smallest S-suffixes starting with each character c such that no L-suffix starts with c, and $suf(\overline{Sy})$ is the other S-suffixes. The important difference with Section 4 is that our new algorithm sorts L- and S-suffixes from sorted S- and L-suffixes, respectively. Obviously, $\mathbf{SA}$ can be computed in $O(N)$ time and in-place in the following steps.

1. Make $\mathbf{SA}_{suf(L)}$ from $\mathbf{SA}_{suf(LMS)}$.
2. Make $\mathbf{SA}_{suf(S)}$ from $\mathbf{SA}_{suf(L)}$.
3. Make $\mathbf{SA}_{suf(\overline{Lx})}$, $\mathbf{SA}_{suf(Lx) \cup suf(Sy)}$, and $\mathbf{SA}_{suf(\overline{Sy})}$ from $\mathbf{SA}_{suf(S)}$.
4. Make $\mathbf{SA}$ by stably merging the aforementioned three arrays considering the first characters as keys.

From Lemma 3, we store $rank(i+1)$ in $\mathbf{B}[j]$ at the same time when storing an induced S-suffix/L-suffix $\mathbf{T}_i$ into $A[j]$ in Step 2/3. Thereafter, the movement of $rank(i+1)$ in $\mathbf{B}$ is synchronized with $\mathbf{T}_i$ in $\mathbf{A}$, namely, $rank(i+1)$ moves to $\mathbf{B}[k]$ when $\mathbf{T}_i$ moves to $\mathbf{A}[k]$. In the end of the algorithm, we have $\mathbf{SA}$ in $\mathbf{A}$ and $\mathbf{\Psi}$ in $\mathbf{B}$. Because each synchronized movement takes $O(1)$ time, the algorithm runs in $O(N)$ time and in-place.

Lemma 3. *In Step 2/3, a pair of $\mathbf{T}_i$ and $rank(i+1)$ can be obtained for all S-suffixes/L-suffixes $\mathbf{T}_i$.*

Proof. In Step 2/3, we read all suffixes $\mathbf{T}_{i+1}$ in reverse-lexicographic/lexicographic order, and store $\mathbf{T}_i$ somewhere in $\mathbf{A}$ if $\mathbf{T}_i$ is an S-suffix/L-suffix. By counting the number of suffixes that have been read, all S-suffixes/L-suffixes $\mathbf{T}_i$ and $rank(i+1)$ can be obtained simultaneously. □

6 Acknowledgement

We wish to thank Takashi Kato and Shunsuke Inenaga for many valuable insights and suggestions, and also Hideo Bannai for the reviews of the draft, and many corrections and advices.

References

[1] Mohamed Ibrahim Abouelhoda, Stefan Kurtz, and Enno Ohlebusch. The enhanced suffix array and its applications to genome analysis. In *Algorithms in Bioinformatics, Second International Workshop, WABI 2002, Rome, Italy, September 17-21, 2002, Proceedings*, pages 449–463, 2002. URL: http://dx.doi.org/10.1007/3-540-45784-4_35, doi:10.1007/3-540-45784-4_35.

[2] Jérémy Barbay, Francisco Claude, Travis Gagie, Gonzalo Navarro, and Yakov Nekrich. Efficient fully-compressed sequence representations. *Algorithmica*, 69(1):232–268, 2014. URL: http://dx.doi.org/10.1007/s00453-012-9726-3, doi:10.1007/s00453-012-9726-3.

[3] M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical report, 1994.

[4] Jingchao Chen. Optimizing stable in-place merging. *Theor. Comput. Sci.*, 302(1-3):191–210, 2003. URL: http://dx.doi.org/10.1016/S0304-3975(02)00775-2, doi:10.1016/S0304-3975(02)00775-2.

[5] Maxime Crochemore and Lucian Ilie. Computing longest previous factor in linear time and applications. *Inf. Process. Lett.*, 106(2):75–80, 2008. URL: http://dx.doi.org/10.1016/j.ipl.2007.10.006, doi:10.1016/j.ipl.2007.10.006.

[6] Johannes Fischer, Volker Heun, and Stefan Kramer. Fast frequent string mining using suffix arrays. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM 2005), 27-30 November 2005, Houston, Texas, USA*, pages 609–612, 2005. URL: http://dx.doi.org/10.1109/ICDM.2005.62, doi:10.1109/ICDM.2005.62.

[7] Gianni Franceschini and S. Muthukrishnan. In-place suffix sorting. In *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings*, pages 533–545, 2007. URL: http://dx.doi.org/10.1007/978-3-540-73420-8_47, doi:10.1007/978-3-540-73420-8_47.

[8] Keisuke Goto and Hideo Bannai. Space efficient linear time Lempel-Ziv factorization for small alphabets. In *Data Compression Conference, DCC 2014, Snowbird, UT, USA, 26-28 March, 2014*, pages 163–172, 2014. URL: http://dx.doi.org/10.1109/DCC.2014.62, doi:10.1109/DCC.2014.62.

[9] Keisuke Goto, Hideo Bannai, Shunsuke Inenaga, and Masayuki Takeda. Fast q-gram mining on SLP compressed strings. *J. Discrete Algorithms*, 18:89–99, 2013. URL: http://dx.doi.org/10.1016/j.jda.2012.07.006, doi:10.1016/j.jda.2012.07.006.

[10] Guy Jacobson. Space-efficient static trees and graphs. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 549–554, 1989. URL: http://dx.doi.org/10.1109/SFCS.1989.63533, doi:10.1109/SFCS.1989.63533.

[11] Juha Kärkkäinen, Peter Sanders, and Stefan Burkhardt. Linear work suffix array construction. *J. ACM*, 53(6):918–936, 2006. URL: http://doi.acm.org/10.1145/1217856.1217858, doi:10.1145/1217856.1217858.

[12] Toru Kasai, Gunho Lee, Hiroki Arimura, Setsuo Arikawa, and Kunsoo Park. Linear-time longest-common-prefix computation in suffix arrays and its applications. In *Combinatorial Pattern Matching, 12th Annual Symposium, CPM 2001 Jerusalem, Israel, July 1-4, 2001 Proceedings*, pages 181–192, 2001. URL: http://link.springer.de/link/service/series/0558/bibs/2089/20890181.htm.

[13] Dong Kyue Kim, Jeong Seop Sim, Heejin Park, and Kunsoo Park. Constructing suffix arrays in linear time. *J. Discrete Algorithms*, 3(2-4):126–142, 2005. URL: http://dx.doi.org/10.1016/j.jda.2004.08.019, doi:10.1016/j.jda.2004.08.019.

[14] Pang Ko and Srinivas Aluru. Space efficient linear time construction of suffix arrays. *J. Discrete Algorithms*, 3(2-4):143–156, 2005. URL: http://dx.doi.org/10.1016/j.jda.2004.08.002, doi:10.1016/j.jda.2004.08.002.

[15] Fugen Li and Gary D. Stormo. Selection of optimal DNA oligos for gene expression arrays. *Bioinformatics*, 17(11):1067–1076, 2001. URL: http://dx.doi.org/10.1093/bioinformatics/17.11.1067, doi:10.1093/bioinformatics/17.11.1067.

[16] Zhize Li, Jian Li, and Hongwei Huo. Optimal in-place suffix sorting. *CoRR*, abs/1610.08305, 2016. URL: http://arxiv.org/abs/1610.08305.

[17] Udi Manber and Eugene W. Myers. Suffix arrays: A new method for on-line string searches. *SIAM J. Comput.*, 22(5):935–948, 1993. URL: http://dx.doi.org/10.1137/0222058, doi:10.1137/0222058.

[18] Giovanni Manzini. Two space saving tricks for linear time LCP array computation. In *Algorithm Theory - SWAT 2004, 9th Scandinavian Workshop on Algorithm Theory, Humlebaek, Denmark, July 8-10, 2004, Proceedings*, pages 372–383, 2004. URL: http://dx.doi.org/10.1007/978-3-540-27810-8_32, doi:10.1007/978-3-540-27810-8_32.

[19] Gonzalo Navarro. A guided tour to approximate string matching. *ACM Comput. Surv.*, 33(1):31–88, 2001. URL: http://doi.acm.org/10.1145/375360.375365, doi:10.1145/375360.375365.

[20] Ge Nong. Practical linear-time $O(1)$-workspace suffix sorting for constant alphabets. *ACM Trans. Inf. Syst.*, 31(3):15, 2013. URL: http://doi.acm.org/10.1145/2493175.2493180, doi:10.1145/2493175.2493180.

[21] Ge Nong, Sen Zhang, and Wai Hong Chan. Two efficient algorithms for linear time suffix array construction. *IEEE Trans. Computers*, 60(10):1471–1484, 2011. URL: http://dx.doi.org/10.1109/TC.2010.188, doi:10.1109/TC.2010.188.